

# Design and Implementation of the PALM-3000 Real-Time Control System

Tuan N. Truong<sup>1</sup>, Antonin H. Bouchez<sup>2</sup>, Rick S. Burruss<sup>1</sup>, Richard G. Dekany<sup>3</sup>,  
Stephen R. Guiwits<sup>3</sup>, Jennifer E. Roberts<sup>1</sup>, Jean C. Shelton<sup>1</sup>, Mitchell Troy<sup>1</sup>

<sup>1</sup>Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA 91109

<sup>2</sup>Giant Magellan Telescope, Pasadena, CA, USA 91109

<sup>3</sup>California Institute of Technology, Pasadena, CA, USA 91125

## ABSTRACT

This paper reflects, from a computational perspective, on the experience gathered in designing and implementing real-time control of the PALM-3000 adaptive optics system currently in operation at the Palomar Observatory. We review the algorithms that serve as functional requirements driving the architecture developed, and describe key design issues and solutions that contributed to the system's low compute-latency. Additionally, we describe an implementation of dense matrix-vector-multiplication for wavefront reconstruction that exceeds 95% of the maximum sustained achievable bandwidth on NVIDIA Geforce 8800GTX GPU.

**Keywords:** PALM-3000, adaptive optics, wavefront processor, dense matrix-vector multiplication, GPU, CUDA

## 1. INTRODUCTION

Adaptive optics (AO) by its real-time nature entails fast response to wavefront changes. For an AO system to work well, commands sent to actuators changing the surface of a deformable mirror (DM) to provide the necessary compensations must be computed while the aberrations are still small. The primary goal in designing the PALM-3000 real-time control (RTC) algorithms is to achieve the highest possible control bandwidth, while assuring stability and robustness in a wide range of atmospheric conditions. From the computational perspective, this implies lowest possible compute-latency. The aim of this paper is to reflect on key design and implementation issues encountered and the solutions developed that contributed to meeting the RTC computational requirements.

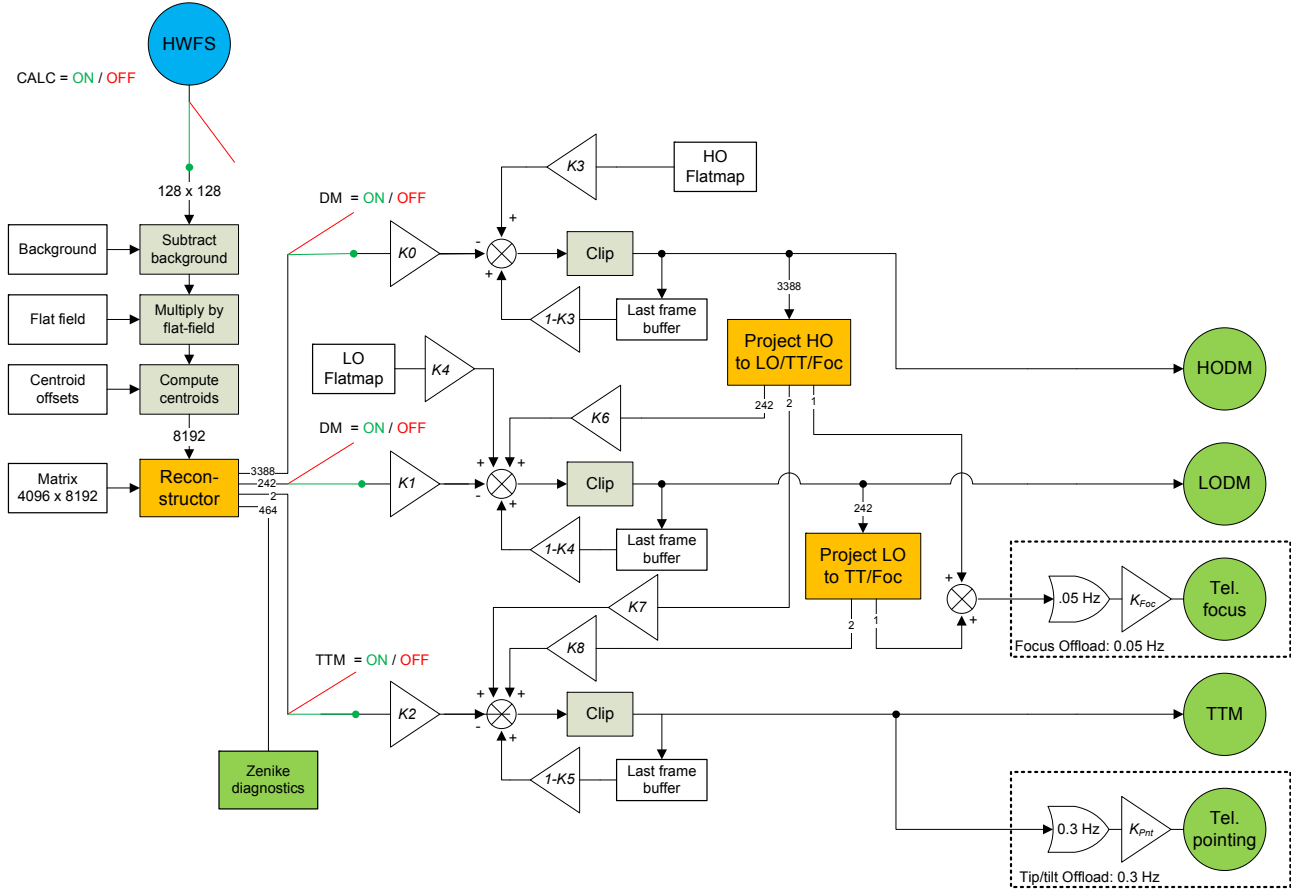
The remainder of the paper is organized as follows. The RTC algorithms that serve as functional requirements driving the architecture developed are first described in Section 2. From there it will be evident that the matrix-vector multiplication (VMM) operation used for wavefront reconstruction is the computational bottleneck. Section 3 presents the distributed multi-GPU architecture. Section 4 provides the latency of each algorithm used. Section 5 describes a GPU implementation of VMM optimized specifically for the RTC reconstructor size on the NVIDIA Compute-Unified-Device (CUDA) architecture<sup>[1]</sup> and compares its performance to a leading method. Section 6 concludes the paper.

## 2. ALGORITHMS

The PALM-3000<sup>[1]</sup> adaptive optics system which has been in operation on the 5.1 meter Hale Telescope at the Palomar Observatory since 2011<sup>[5]</sup>, employs a 64x64 subaperture Shack-Hartmann high-order wavefront sensor (HOWFS) for measurement, and three adaptive mirrors for correction. The mirrors include a tip/tilt mirror (TTM), a low-order deformable mirror (LODM) of 241 active actuators, and a high-order deformable mirror (HODM) of 3388 active actuators. The two DMs are operated in a woofer-tweeter configuration, with the LODM providing large stroke at low spatial frequencies ( $\pm 2.5 \mu\text{m}$  surface stroke with 17 actuators across the telescope pupil) and the HODM providing fine correction at high spatial frequencies ( $\pm 0.5 \mu\text{m}$  surface stroke with 64 actuators across the pupil).

Figure 1 depicts end-to-end dataflow of the RTC with all computation steps used to convert pixel values from the

HOWFS to commands issued to the DMs. As shown, the raw pixels are read out in frames of 128x128 unsigned 16-bit values, background-subtracted and multiplied by the flat field before inputting to the current active centroiding algorithm to obtain a vector of 8192 slope measurements or centroids. Background subtraction, which compensates for the camera bias, sky background, and other additive flux, involves simply subtracting a background, or dark, value acquired with no guide star light from each raw pixel value. Flat-field correction, which compensates for individual pixel gain or QE variations in the wavefront sensor camera, requires only a multiplication operation per pixel.



**Figure 1: PALM-300 Servo Control**

The centroid vector is computed with a weighted sum over the flat-fielded pixels, divided by the total light in the subaperture. Equations below show the quad-cell centroiding algorithm used in the RTC. Therein  $c$  denotes the centroid vector output,  $d$  the centroid offset vector,  $f_{\min}$  the user-selectable minimum flux,  $g$  the weights  $+1, +1, -1, -1$  and  $+1, -1, +1, -1$  for resolving  $x$  and  $y$  slopes,  $s$  the pixel values from flat fielding,  $i$  the index of subaperture,  $j$  the index of pixels within a subaperture, and  $N$  the total number of subapertures.

Note the clamping of the denominator to  $f_{\min}$  rather than zero. The effect of this is to retain AO performance for the faintest objects rather than simply quitting. This centroiding technique known as *clamped denominator*<sup>[6]</sup> along with reconstructor upgrades<sup>[7]</sup> have shown performance improvements on sky when operating PALM-3000 on faint targets.

$$\begin{aligned}
\mathbf{S}_i &= \sum_{j=1}^N \mathbf{s}_{i,j} \\
\mathbf{c}_i &= \mathbf{w}_i \left( \sum_{j=1}^N \mathbf{g}_{i,j} \mathbf{s}_{i,j} - d_i \mathbf{S}_i \right) \\
\text{if } \mathbf{S}_i > f_{\min} \text{ then } \mathbf{w}_i &= 1 / \mathbf{S}_i \text{ else } \mathbf{w}_i = 1 / f_{\min}
\end{aligned}$$

Next, the output centroid vector containing 8192 slope measurements (2 per subaperture) is multiplied with the current active reconstructor matrix to generate 4096 residuals that estimate the wavefront errors. Of the 4096 residuals, only 3388 values go to the HODM servo loop, 241 to the LODM, and 2 to the TTM, to generate the corresponding DM and TTM commands, with the remaining values available for diagnostics information.

Computing the DM actuator command requires the previous timestep actuator command vector, the flatmap voltages to which we want the DM to relax in absence of an error signal, the current residual vector from the reconstructor, and two servo gain coefficients. Computing a LODM actuator command requires an additional servo gain coefficient and the previous timestep LODM offload vector, which is the projection of the HODM onto the LODM calculated as the product of the HODM-to-LODM projection matrix and the HODM actuator command.

Computing the TTM command requires four servo coefficients, the current residual value, the previous timestep actuator command, the previous timestep LODM offload vector, and the previous TTM offload vector which represents the projection of the LODM onto the TTM which is calculated by multiplying the LODM-to-TTM projection matrix with a LODM actuator command.

The HODM-to-LODM and LODM-to-TTM projections are combined into one VMM operation of size 256x4096 with the estimated wavefront error vector. Note since this operation is done after the actuator commands have been issued, its execution time is excluded from the total compute-latency of the system.

Camera mode	Pupil Sampling	Available subapertures $N_{sub}$	Illuminated subapertures (>50%)	
			No obscuration	Central obscuration
0	64	4096	3145	2800
1-3, 9	32	1024	792	704
4-6, 10-11	16	256	200	176
7-8, 12	8	64	52	48

**Table 1:** Number of subapertures for various observing modes

There are optimization techniques<sup>[8]</sup> that can be performed independent of architecture. First, to minimize idle time, we overlap communication with computation, and ensure that the work done is evenly distributed across the system. Specifically, transfers of pixels are done via DMA in blocks of half-frames. This permits calculations to start as soon as a block has been received and progress while the second half of pixels making up the frame is read out. Section 4 confirms the latency improvement expected of half-frame blocking. This technique requires one buffer for calculations and at least another for transfers. However, we allocated a ring of buffers for transfer to tolerate occasional jitter due to process preemption by kernel activities on standard Linux. Calculations should therefore be fast enough to catch up quickly once the process is unblocked. Second, we exploit non-blocking communication. That is, we issue multiple asynchronous sends and receives to allow simultaneous transmission and reception of data from multiple sources to multiple destinations. To ensure load balanced, equal but different portions of data are calculated and transferred using the same number of instructions. Toward this end, we divide and coalesce memory fragments, one from each data

product, so that every processing node makes the same minimum number of equal size DMA transfers. This data shuffling in memory is justified since memory bandwidth is orders of magnitude faster than communication bandwidth.

Contrary to our stated goal of minimizing latency, the current version of the software receives and processes in half-frames of 64x128 pixels each, regardless of the HOWFS camera mode currently operating. It also multiplies the reconstructor with all subapertures whether illuminated or not. Thus the vector of measured slopes *output* by RTC centroiding algorithms always contain 8192 elements, in spite of various observing modes with different number of subapertures (see Table 1). Having said that, the number of elements actually *computed* does vary according to observing modes. VMM is also currently done on dense 4096x8192 matrices, despite the total number of controlled actuators, not including those slaved electrically is only 3631 (3388 for HODM, 241 for LODM and 2 for TTM). Future versions will correct these issues if funding permits.

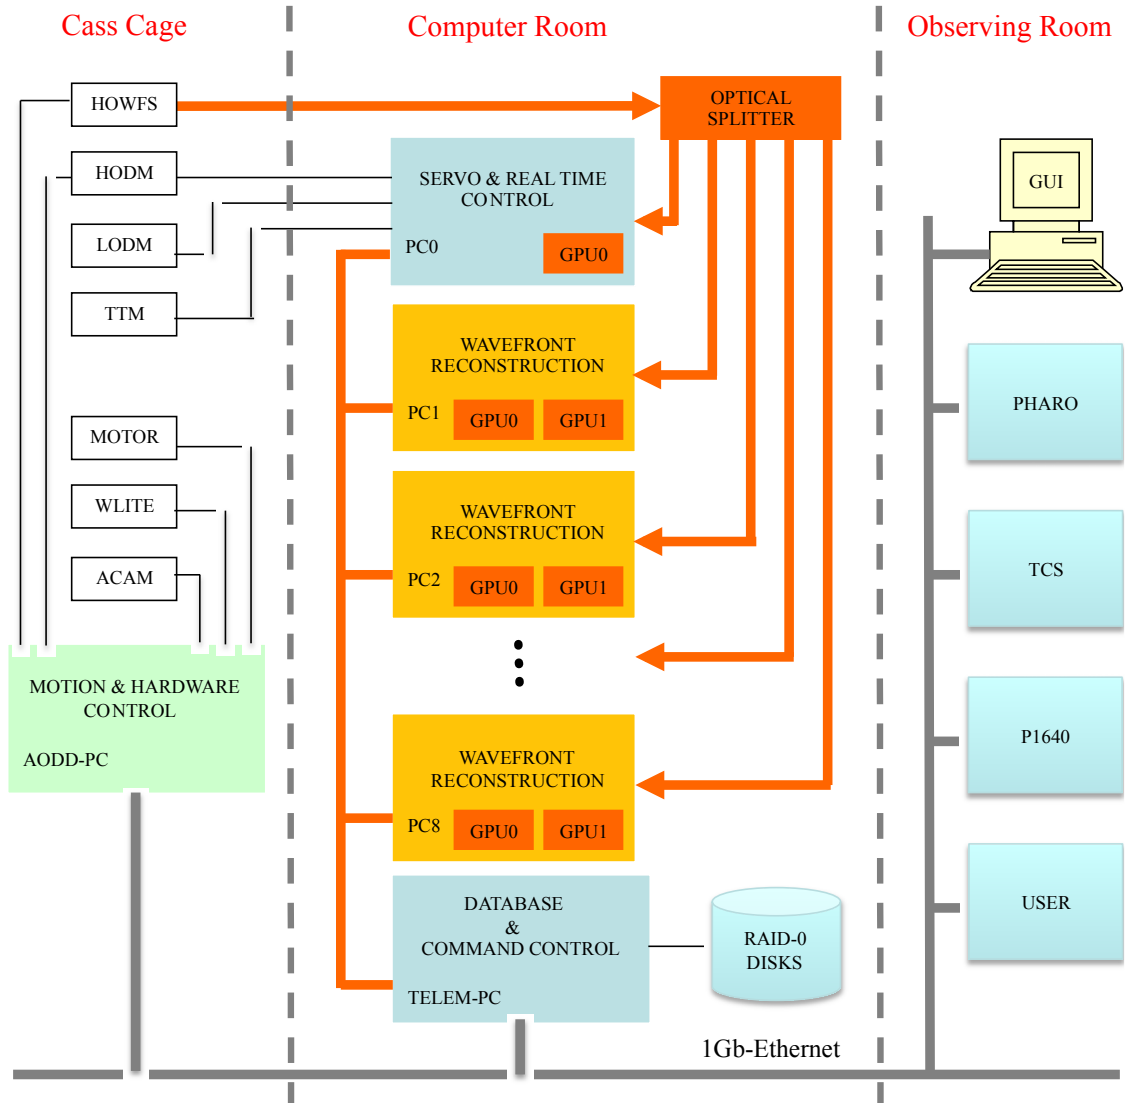
### 3. ARCHITECTURE

This section presents the architecture of the PALM-3000 computer system and describes how the RTC algorithms are mapped and parallelized on the hardware adopted. The RTC architecture is shown in Figure 2 along with external entities that use it. The system consists of five major functional subsystems, RTC, AODD, AOCA, AODB, and GUI, that spread over three separate locations, the cass-cage, the computer room, and the users/observing area, that in turn interconnect via 1Gbit Ethernet. Non-critical or low speed data use this network, while real-time high-speed data go on the dedicated fiber link with direct point-to-point connections to the cass cage or via the Quadric switch.

The cass cage holds the optics, DMs, TTM, HOWFS and all hardware that must absolutely be in proximity for proper operation such as motors, fiber optic receivers/transmitters, acquisition camera and a Linux PC that hosts the necessary device drivers. The observing area, on the other end, represents the logical location of the users, which may be either operators and scientists that control the system via the graphical user interface software provided or external systems that interface with PALM-3000 such as TCS, Pharo or P1640.

Shown in the center of the figure is the computer room that houses ten AO system PCs. Nine computers named *pc0* to *pc8* run the RTC algorithms described, while the 10<sup>th</sup> named *telem* hosts the database (AODB) and command/automation (AOCA) subsystems. All ten PCs are HP xw9400 running the standard RedHat Linux 2.6.18 kernel with two dual-core processors AMD Opteron 2220 2.8GHz, 128KB L1 cache and 1MB L2 cache *per core*, 4GB memory DDR2 667MHz, and two NVIDIA GeForce 8800GX GPUs of full PCIe x16 bandwidth. Exceptions are *pc0*, which has only one GPU, and *telem*, which has 16GB of memory but no GPUs. All PCs are linked via an ultra-low latency high bandwidth Quadrics switch that delivers over 900 MB/s of user-space-to-user-space bandwidth each direction for a total of 14.4 GB/s of bisectional bandwidth and latency of only 1.8 $\mu$ s.

Pixel values from the HOWFS camera are replicated by the optical splitter and sent to 8 processing nodes (*pc1* to *pc8*) as they are read out. Once one half of the pixels have been received, each GPU computes the same centroids and multiplies by a different portion of the reconstruction matrix to obtain a “half-computed” wavefront error (residual) vector for one-sixteenth of the output actuators. When this process repeated on the second half of the pixels making up the frame completes, the two half-computed residual vectors are summed and then transmitted to the central node (*pc0*). Once all residuals have been received by *pc0*, the servo loop begins. The servo algorithm computes the output actuator command, sends it off to the respective mirror, and then calculates the offload projections before repeating the loop on the next residual vector.



**Figure 2:** PALM-3000 hardware architecture. HOWFS denotes the high-order wavefront sensor, HODM the high-order deformable mirror, LODM the low-order deformable mirror, TTM the tip/tilt mirror, MOTOR all motors controlled, WLITE the white light source, ACAM the acquisition camera, AODD the device drivers, AOCA the command and automation software, TELEM the PC hosting the database, AOCA and RAID-0 file system, GUI the graphical user interface provided to control the PALM-3000 system, PHARO the Palomar high angular resolution observer, a near-infrared camera built by Cornell University, TCS the telescope control system, P1640 the advanced coronagraph and high-spectral imaging system built by American Museum of Natural History, USER any process/system subscribing to PALM-3000 data.

*pc0* additionally synchronizes and manages all traffic in and out of the eight slaves. All messages are tagged with a 32-bit control word containing a session number and a sequence number that uniquely identifies a message in that session. Sequence number 0 is reserved for messages containing commands and responses. Each slave increments its own sequence number every time it sends out a new message containing data. *Pc0* uses the sequence numbers to identify and coalesce all residuals belonging to the same frame, provided that the data messages have the same current session number. Otherwise, they are dropped which could happen when one slave detects a frame error and notifies *pc0* quicker than others. The session number, on the other hand, is incremented by *pc0* and broadcasted to the slaves each time all synchronize in a rendezvous. Slaves synchronize in a rendezvous only when commanded by *pc0*. Otherwise they run freely, even right after they have detected a hardware fault. All user commands from AOCA go through *pc0* which

parses them out to the slaves, waits for the results from each, and replies on their behalf before synchronizing in a rendezvous with all of them.


All nine RTC computers log data directly to the RAID-0 system located on *telem* via AODB, a publish-subscribe software layer built in-house from the ground up on top of Berkeley DB<sup>[9]</sup> and used by all PALM-3000 software components for communication and database services. To enable seamless distributed processing across machines, we use bidirectional proxy components that are interposed between user components and remote provider components.

The RTC software was inspired by the split-phase, event-driven concurrency model of TinyOS<sup>[10]</sup>. To achieve non-blocking as dictated by our optimization goal stated earlier, all long latency operations are performed in split-phase, meaning the operations are implemented using two functions, one called *command* which is the request and which returns immediately, and the other called *event* which is signaled upon completion. Although an execution may be rooted in an event (typically representing a hardware interrupt), more often than not, it is rooted in what are called tasks, which have no strict timing requirements and are scheduled by a scheduler to run at a later time. Tasks are a deferred computation mechanism. They run to completion and unlike events, they do not preempt one another. Events can however preempt them. Any running codes may *post* (schedule) a task for later execution. For almost all of the PALM-3000 software, this split-phase, event-driven execution model suffices. But to drive multiple GPUs and to avail of the parallelism afforded by multiple CPU cores, threads are required. So we extended nesC<sup>[11, 12]</sup>, the language of PALM-3000 software, with thread support by enhancing its task and event-signaling model. Specifically, we introduced a mechanism for associating a nesC task with a thread of execution, which may be an existing thread or one newly created by the task scheduler. We use threads like we use external hardware. In this manner, threads signaling events are akin to hardware issuing interrupts. Hence, events signaled and commands invoked by tasks associated with threads must be declared asynchronous (i.e., *async* in nesC). We also enhanced the language to allow applications to post a gang of tasks, where any number of them could be defined outside the current local scope (i.e., in other components, modules or processes). We use this feature to synchronize and notify components of system-wide events. One such event is application shutdown, which entails closing of all opened Berkeley DB handles. The same mechanism associating a task with a thread of execution is also used to specify any task's priority, including a pure nesC task. We chose this approach to retain the static race-condition analysis feature of the compiler. A problematic alternative would be having applications deal directly with a thread library.

Like TinyOS, the RTC software is built out of components, which are assembled according to a wiring specification to form an application. That wiring specification, called *interface* in nesC, is the only access point to components. We designed interfaces to contain only the minimum functionality required for most usage scenarios. Functions that enable reconfiguration and customization are not part of the interface. Instead they are enabled at runtime via component-specific command lines switches, environment variables, and configuration files. For example, on our system, a component using the interface to subscribe to remote data deals with only one function to receive the data. It does not make any calls to initialize the providing component that implements the interface, change the subscribed data rate, or specify programmatically the whereabouts of the data and its server. The interface simply does not provide any of those functions. This design approach has resulted in not only a system that is more extensible but also components that are simpler to develop.

## 4. PERFORMANCE

Figure 3 lists the execution times in microseconds, averaged over 1 million invocations ignoring the first, of each algorithm computed by the RTC as described previously. It shows clearly the benefit of half-frame blocking over the obvious full-frame. Measuring from the time the last pixel was received by the application up to the instant actual commands are ready to be issued to the mirrors, the total elapsed time is 205 $\mu$ s for half-frame blocking versus 281 $\mu$ s for full-frame. Not included in these totals nor shown in the figure are the latencies of interactor check, the DM driver and the offload calculation that follows immediately in parallel. The actuator check requires 35 $\mu$ s with room for improvement. The DM driver's latency currently at 125 $\mu$ s is expected to drop down to 25 $\mu$ s later this summer. The offload latency of 80 $\mu$ s, which included 13 $\mu$ s for the launch overhead, is excluded from the total compute-latency because the calculation starts after the commands have been issued. From the transfer times shown, it is clear that getting data in and out of GPUs is not of an issue, provided that the input and output buffers were allocated from non-pageable memory.



**Figure 3:** RTC latencies in microseconds. Shown is the processing timeline of the RTC servo algorithm for two frame periods at 2 KHz frame rate. Measuring from the last pixel received by the application up to the point actuator commands are ready to be issued to the mirror, the total latency is 205μs for half-frame blocking and 281μs for full-frame. Not included in these totals nor shown in the figure are the latencies of interactuator check and DM driver. “Upload All” denotes uploading (GPU-to-host) of all data, which include pixels, flux, centroids, and residuals. The transfer size equals the total divided by the number of GPUs in the system.

## 5. MATRIX VECTOR MULPLICATION

This section presents an improved VMM algorithm used by the RTC. Figure 4 shows the core loop of the code on the NVIDIA CUDA architecture. Table 4 compares its performance with the leading method of Fujimoto<sup>[13]</sup>. The times shown were averaged over 1 million kernel invocations ignoring the first, for several matrix sizes of interest to PALM-3000, excluding data transfers between CPU and GPU and kernel launch overhead. The thread blocks used for the comparison are 16x4 for the PALM-3000 method and 16x16 as suggested for [13].

Matrix Size	Fujimoto	Palm-3000	Speedup
256 x 4096	75.5	61.0	1.23
256 x 2048	40.0	32.0	1.25
512 x 4096	150.0	129.0	1.16
512 x 2048	92.5	67.5	1.37

**Table 4:** Matrix-vector multiplication performance comparison. Speedup is the ratio of two times (μs).

As shown, the performance of the RTC method on small matrices is unmatched and due to the superior caching design of the GPU texture memory. For a bandwidth-bounded algorithm like VMM, the execution time can be predicted by the equation.

$$\text{Time} = \text{Kernel launch overhead} + \frac{\text{Bandwidth required}}{72 \text{ GB/s}}$$

Using our hardware/software setup, the maximum sustained achievable bandwidth reported by *bandwidth* (an application provided in the NVIDIA software development toolkit) is 72GB/s, and the launch overhead 4 $\mu$ s if the kernel is launched right after another and 13 $\mu$ s otherwise. For the matrix size of 256x4096 operated on by each GPU, the RTC algorithm attains over 95% of the maximum achievable memory bandwidth of the GeForce 8800GTX GPU, providing speedup of 1.23x over [13]. However, for large matrices such as 2048x4096 and 4096x4096 that are not of interest, the RTC method is outrun by 10% to 15% respectively.

```
// Y = AX
// A : matrix cached in texture memory referenced via textureA.
// X : vector cached in texture memory referenced via textureX.
//
// This code is illustrative.
// Assumed to be launched with RxC threads per block,
// and the number of thread blocks equal to heightA/R.
//
// tx, ty : thread id.
// bx : block id.
//
__shared__ float y[R][C];

row = bx*R + ty;
y[ty][tx] = 0;

for (i = 0; i < widthA/4;
     i+= C)
{
    float4 a = tex2D( textureA, tx+i, row );
    float4 x = tex1Dfetch( textureX, tx+i );
    y[ty][tx] += a.x*x.x + a.y*x.y + a.z*x.z + a.w*x.w;
}
__syncthreads();

// Sum y[ty][*] and store result in Y[row]
// in binary-tree fashion with warp specialization.
'
```

**Figure 4:** Partial kernel codes for matrix-vector multiplication in RTC

## 6. CONCLUSIONS

We described the algorithms used for real time control of the PALM-3000 adaptive optics system and presented a scalable architecture that satisfies its computational requirements. The key components we pointed out that contributed to the system's low latency were the optical splitter, the GPUs, and the Quadrics switch. With larger interconnects and additional if not faster GPUs, we expect the architecture could easily support much larger AO systems at higher rates and lower VMM latency. We also presented an improved algorithm for matrix-vector multiplication on the NVIDIA CUDA architecture that reaches within 5% of achievable sustained GPU bandwidth and delivers speedup ranging from 1.16x to 1.37x over a leading method.

## ACKNOWLEDGEMENTS

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the California Institute of Technology and the National Aeronautics and Space Administration..



## REFERENCES

- [1] NVIDIA, "CUDA Programming Guide", Version 1.0, June 23, 2007.
- [2] Dekany, R.G., et al., "PALM-3000: visible light AO on the 5.1 meter telescope", Proc. SPIE Vol. 6272, 2006.
- [3] Bouchez, B.H., et al., "The PALM-3000 High-Order Adaptive Optics Upgrade", Proc. SPIE Astronomical Instrumentation 2008
- [4] Truong, T.N., et al., "Real-Time Wavefront Control for the PALM-3000 High Order Adaptive Optics", Proc. SPIE Astronomical Instrumentation 2008.
- [5] Dekany, R.G., et al., in preparation
- [6] Shelton, C., "Denominator-Free Control Algorithms for Adaptive Optics", Proc. SPIE Vol. 3126, p. 455-459, 1997
- [7] Shelton, C., Troy M., Bouchez, A., Roberts, J., Trinh, T., Truong, T., "NGS and LGS Adaptive Optics – Improving Faint Light Performance", Center For Adaptive Optics Fall Retreat 2009.
- [8] Becker, A., Venkataraman, R., Kale, L., "Patterns for Overlapping Communication and Computation", Dept. of Computer Science, Univ. of Illinois at Urbana-Champaign, Urbana, IL.
- [9] Olson, M.A., et al., "Berkeley DB", Proc. of the FREENIX Track, 1999 USENIX Annual Technical Conference.
- [10] Levis, P., "TinyOS Programming", <http://www.tinyos.net>, October 27, 2006.
- [11] Gay, D., et al., "The nesC language: A holistic approach to networked embedded systems", Proceedings of the ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation, p. 1-11, 2003.
- [12] Gay, D., et al., "nesC 1.2 Reference Manual", August 2005.
- [13] Fujimoto, N., "Dense Matrix-Vector Multiplication on the CUDA Architecture", Parallel Processing Letters 2008, Vol. 18, No. 4 p. 511-530.